

# Meaningful Clustered Forest: an Automatic and Robust Clustering Algorithm

Mariano Tepper

*Departamento de Computación, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires*  
e-mail: [mtepper@dc.uba.ar](mailto:mtepper@dc.uba.ar)

Pablo Musé

*Instituto de Ingeniería Eléctrica, Facultad de Ingeniería, Universidad de la República,*  
e-mail: [pmuse@fing.edu.uy](mailto:pmuse@fing.edu.uy)

Andrés Almansa

*CNRS - LTCI UMR5141, Telecom ParisTech*  
e-mail: [andres.almansa@telecom-paristech.fr](mailto:andres.almansa@telecom-paristech.fr)

**Abstract:** We propose a new clustering technique that can be regarded as a numerical method to compute the proximity gestalt. The method analyzes edge length statistics in the MST of the dataset and provides an a contrario cluster detection criterion. The approach is fully parametric on the chosen distance and can detect arbitrarily shaped clusters. The method is also automatic, in the sense that only a single parameter is left to the user. This parameter has an intuitive interpretation as it controls the expected number of false detections. We show that the iterative application of our method can (1) provide robustness to noise and (2) solve a masking phenomenon in which a highly populated and salient cluster dominates the scene and inhibits the detection of less-populated, but still salient, clusters.

**Keywords and phrases:** clustering, minimum spanning tree, a contrario detection.

## 1. Introduction

Clustering is an unsupervised learning method that seeks to group observations into subsets (called clusters) so that, in some sense, intra-cluster observations are more similar than inter-cluster ones. Despite its intuitive simplicity, there is no general agreement on the definition of a cluster. In part this is due to the fact that the notion of cluster cannot be trivially separated from the context. Consequently, in practice different authors provide different definitions, usually derived from the algorithm being used, rather than the opposite.

Unfortunately, the lack of a unified definition makes it difficult to find a unifying clustering theory. A plethora of methods to assess or classify clustering algorithms have been developed, some of them with very interesting results. To

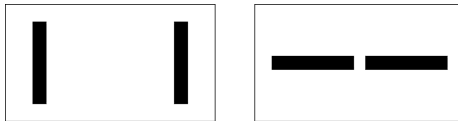


FIG 1. *Two experiments with black bars. We perceive the bars on the left as more separated than the ones on the right. Nevertheless, there exists distances between sets that cannot capture the difference.*

cite a few [7, 20, 21]. For a broad perspective on clustering techniques, we refer the reader to the excellent overview of clustering methods recently reported by Jain [18].

Human perception is extremely adapted to group similar visual objects. Based on psychophysical experiments using simple 2D figures, the Gestalt school studied the perceptual organization, and identified a set of rules that govern human perception [25]. Each of these rules focuses on a single quality, or gestalt, many of which have been unveiled over the years.

One of the earlier and most powerful gestalts is proximity, which states that spatial or temporal proximity of elements may be perceived as a single group. Of course, the notion of distance is heavily embedded in the proximity gestalt. This is clearly illustrated in Figure 1. Two possible distances between the bars  $B_1$  and  $B_2$  that could be considered are

$$d_M(B_1, B_2) = \max_{\substack{p_1 \in B_1 \\ p_2 \in B_2}} \|p_1 - p_2\|,$$

$$d_m(B_1, B_2) = \min_{\substack{p_1 \in B_1 \\ p_2 \in B_2}} \|p_1 - p_2\|.$$

In this particular example  $\|\cdot\|$  denotes the euclidean norm. According to distance  $d_M$ , the bars are exactly at the same distance in both experiments, while according to distance  $d_m$  the bars on the right are closer to each other. In this case, the distance  $d_m$  seems to be more consistent with our perception.

The conceptual grounds on which our work is based were laid by Zahn in a seminal paper from 1971 [26]. Zahn faced the problem of finding perceptual clusters according to the proximity gestalt and proposed three key arguments:

1. **Only inter-point distances matter.** This imposes graphs as the only suitable underlying structure for clustering.
2. **No random steps.** Results must remain stable for all runs of the detection process. In particular, random initializations are forbidden.
3. **Independence from the exploration strategy.** The order in which points are analyzed must not affect the outcome of the algorithm.

These conceptual statements, together with the preference for  $d_m$  over  $d_M$  or other distances between sets, led Zahn to use the Minimum Spanning Tree (MST) as a building block for clustering algorithms. (The MST is the tree structure induced by the distance  $d_m$  [9].) Recently, psychophysical experiments

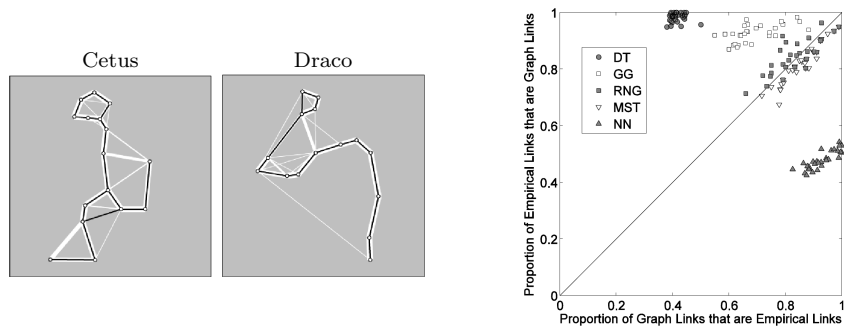


FIG 2. Left and middle: example constellations shown in black and the aggregated empirical structure shown in white. The number of persons that chose an edge is represented by the edge's width. Right: proportional overlap between graph and empirical structure links for Delaunay Triangulation (DT), Gabriel Graph (GG), Relative Neighborhood graph (RNG), Minimum Spanning Tree (MST), and Nearest Neighbors (NN). Each data point represents one of the 30 stimuli. Reproduced from [11].

performed by Dry et al. [11] supported this choice. In these experiments individuals were asked to connect points of 30 major star constellations, to show the structure they perceive. Two examples of constellations are shown in Figure 2. The outcome of these experiments was that, among five relational geometry structures, the MST and the Relative Neighborhood Graph (RNG) exhibit the highest degree of agreement with the empirical edges. In the RNG, two points  $p$  and  $q$  are connected by an edge whenever there does not exist a third point  $r$  that is closer to both  $p$  and  $q$  than they are to each other. The MST is a subgraph of the RNG. Nonetheless the diagonal variance of both groups might suggest that sometimes other links not present nor in the MST nor in the RNG are used.

From a theoretical point of view, Carlsson and Mémoli [7] proved very recently that the single-link hierarchy (i.e. a hierarchical structure built from the MST, as explained later) is the most stable choice and has good convergence properties among other hierarchical techniques.

Zahn [26] suggested to cluster a feature set by eliminating the inconsistent edges in the minimum spanning tree. That is, instead of constructing a MST and as a consequence of the eliminations, a minimum spanning forest is built.

Since then, variations of the limited neighborhood set approaches have been extensively explored. The criteria in most works are based on local properties of the graph. Since perceptual grouping implies an assessment of local properties versus global properties, exclusively local methods must be discarded or patched. For example, Felzenszwalb and Huttenlocher [12] and Bandyopadhyay [3] make use of the MST and RNG respectively. However, in order to correct local observations and to produce a reasonable clustering, they are forced to consider additional ad hoc global criteria.

The computation of the MST requires previous computation of the complete graph. This is a major disadvantage of MST-based clustering methods, that

impose severe restrictions both on time and memory. The obvious workaround is to prune a priori the complete graph (e.g. in image segmentation, the image connexity is exploited), but unfortunately it might produce artifacts in the final solution. In a recent work Tepper et al. [24] proposed an efficient method to compute the MST on metric datasets. The use of this method allows for a significant performance boost over previous MST-based methods (e.g. [5, 12]), thus permitting to treat large datasets.

From an algorithmic point of view, the main problem with the Gestalt rules is their qualitative nature. Desolneux et al. developed a detection theory which seeks to provide a quantitative assessment of gestalts [10]. This theory is often referred as Computational Gestalt Theory and it has been successfully applied to numerous gestalts and detection problems [6, 15, 22]. It is primarily based on the Helmholtz principle which states that no structure is perceived in white noise. In this approach, there is no need to characterize the elements one wishes to detect but contrarily, the elements one wishes to avoid detecting.

In the light of this framework, Desolneux et al. analyzed the proximity gestalt, proposing a clustering algorithm [10]. It is founded on the idea that clusters are groups of points contained in a relatively small area. In other words, by counting points and computing the area that encloses them, one can assess the exceptionality of a given group of points.

The method proposed by Desolneux et al. [10] suffers from some problems. First, it can only be applied to points in an Euclidean 2D space. Second, in order to compute the enclosing areas, the space has to be discretized a priori and such discretization is used to compute the enclosing areas; of course, different discretizations lead to different results. Last, two phenomena called collateral elimination and faulty union in [5] occur when an extremely exceptional cluster hides or masks other less but still exceptional ones.

Cao et al. [5] continued this line of research extending the clustering algorithm to higher dimensions and corrected the collateral elimination and faulty union issues, by introducing what they called indivisibility criterion. However, as their method is also based on counting points on a given region, it is still required that a set of candidate regions is given a priori. The set of test regions is defined to be a set  $\mathcal{R}$  of hyper-rectangles parallel to the axes and of different edge lengths, centered at each data point. The choice of  $\mathcal{R}$  is application specific since it is intrinsically related to cluster size/scale. For example, an exponential choice for the discretization of the rectangle space is made by Cao et al. [5] introducing a bias for small rectangles (since they are more densely sampled). Then each cluster must be fitted by an axis-aligned hyper-rectangle  $R \in \mathcal{R}$ , meaning that clusters with arbitrary shapes are not detected. Even hyper-rectangular but diagonal clusters may be missed or oversegmented. A probability law modeling the number of points that fall in each hyper-rectangle  $R \in \mathcal{R}$ , assuming no specific structure in the data, must be known a priori or estimated. Obviously, this probability depends on the dimension of the space to be clustered.

Recently Tepper et al. [23] introduced the concept of graph-based a contrario clustering. A key element in this method is that the area can be computed from a weighted graph, where the edge weight represents the distance between two

a point must be similar	
to all points in the cluster	to at least one point in the cluster
$k$ -means	single-link algorithm [13]
Cao et al. [5]	Mean Shift [8]
Tepper et al. [23]	Felzenszwalb and Huttenlocher [12]

TABLE 1

*Conceptually there are two different ways to form a cluster. To belong to a cluster a point must be similar to all points in the cluster or to at least one point in the cluster. All algorithms explicitly or implicitly chose one or the other.*

points, using non-parametric density estimation. Since only distances are used, the dimensionality of the problem is reduced to one. However, since this method is conceived for complete graphs, it suffers from a high computational burden.

There is an additional concept behind clustering algorithms that was not stated before: a point, to belong to a cluster, must be similar to all points in the cluster or only to some of them? All the described region-based solutions imply choosing the first option since, in some sense, all distances within a group are inspected. Table 1 shows on which side some algorithms are. Since our goal is to detect arbitrarily shaped clusters, we must place ourselves in the second group. We can do this by using the MST.

Our goal is to design a clustering method that can be considered a quantitative assessment of the proximity gestalt. Hence in Section 2 we propose a clustering method based on analyzing the distribution of distances of MST edges. The formulation naturally allows to detect clusters of arbitrary shapes. The use of trees, as minimally connected graphs, also leads to a fast algorithm. The approach is fully automatic in the sense that the user input only relates to the nature of the problem to be treated and not the clustering algorithm itself. Strictly speaking it involves one single parameter that controls the degree of reliability of the detected clusters. However, these methods can be considered parameter-free, as the result is not sensitive to the parameter value. Results on synthetic but challenging sets are presented in Section 3.

As the method relies on the sole characterization of non-clustered data, it is thus capable of detecting non-clustered data as such. In other words, in the absence of clustered data, the algorithm yields no detections. In Section 4 it is also shown that, by iteratively applying our method to datasets with clustered and unclustered data, it is possible to automatically separate both classes.

In Section 5, we finally illustrate a masking phenomena where a highly populated cluster might occlude or mask less populated ones, showing that the iterative application of the MST-based clustering method is able to cope with this issue, thus solving very complicated clustering problems.

Results on three-dimensional examples of the complete process are presented in Section 6 and we expose some final remarks in Section 7.

## 2. A New Clustering Method: Proximal Meaningful Forest

We now propose a new method to find clusters in graphs that is independent from their shape and from their dimension. We first build a weighted undirected graph  $G = (X, E)$  where  $X$  is a set of features in a metric space  $(M, d)$  and the weighting function  $\omega$  is defined in terms of the corresponding distance function

$$\omega((v_i, v_j)) = d(x_i, x_j). \quad (1)$$

### 2.1. The Minimum Spanning Tree

Informally, the Minimum Spanning Tree (MST) of an undirected weighted graph is the tree that covers all vertices with minimum total edge cost.

Given a metric space  $(M, d)$  and feature set  $X \subseteq M$ , we denote by  $G = (X, E)$  the undirected weighted graph where  $E = X \times X$  and the graph's weighting function  $\omega : E \rightarrow \mathbb{R}$  is defined as

$$\omega((x_i, x_j)) = d(x_i, x_j) \quad \forall x_i, x_j \in X. \quad (2)$$

The MST  $T = (X, E_T)$  of the feature set  $X$  is defined as the MST of  $G$ . A very important and classical property of the MST is that a hierarchy of point groups can be constructed from it.

**Notation 1.** Let  $T = (X, E_T)$  be the minimum spanning tree of  $X$ . For a group of points  $C \in X$ , we denote

$$E(C) = \{(v_i, v_j) \mid v_i, v_j \in C \wedge (v_i, v_j) \in E_T\} \quad (3)$$

The edges in  $E(C)$  are sorted in non-decreasing order, i.e.

$$\forall e_i, e_j \in E(C), \quad i < j \Rightarrow \omega(e_i) \leq \omega(e_j)$$

**Definition 1.** Let  $T = (X, E_T)$  be the minimum spanning tree of  $X$ . A component  $C \subseteq X$  is a set such that the graph  $G = (C, E(C))$  is connected and

- $\exists v \in V, C = \{v\}$  or
- $\nexists C' \in X, C \subset C' \quad \wedge \quad \omega_{\max}(C) > \omega_{\max}(C'),$

where  $\omega_{\max}(C) = \max_{e \in E(C)} \omega(e)$ . A single-link hierarchy  $\mathcal{T}$  is the set of all possible components.

It is important to notice what the single-link hierarchy implies: given two components  $C_1, C_2 \in \mathcal{T}$ , it suffices that there exists a pair of vertices, one in  $C_1$  and one in  $C_2$  that are sufficiently near each other to generate a new component  $C_F \in \mathcal{T}$ , such that  $C = C_1 \cup C_2$  and

$$\omega_{\max}(C_F) = \min_{\substack{v_i \in C_1, v_j \in C_2 \\ (v_i, v_j) \in E_T}} \omega((v_i, v_j)). \quad (4)$$

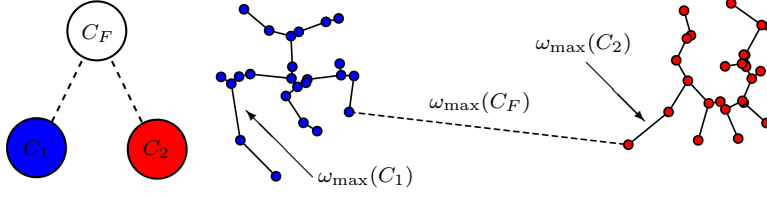


FIG 3. Part of a minimal spanning tree. The blue node set and the red node set are linked by the dashed edge, creating a new node in the minimal spanning tree.

An example is depicted in Figure 3. The direct consequence of this fact is that the use of the single-link hierarchy for clustering provides a natural way to deal with clusters of different shapes.

All minimum spanning tree algorithms are greedy. From Definition 1 and Equation 4, in the single-link hierarchy the component  $C_F = C_1 \cup C_2$  is the father of  $C_1$  and  $C_2$  and

$$\omega_{\max}(C_F) \geq \omega_{\max}(C_1) \quad (5)$$

$$\omega_{\max}(C_F) \geq \omega_{\max}(C_2). \quad (6)$$

With the objective of finding a suitable partition and to the best of our knowledge, Felzenszwalb and Huttenlocher [12] were the first to compare  $\omega_{\max}(C_F)$  with  $\omega_{\max}(C_1)$  and  $\omega_{\max}(C_2)$ , with an additional correction factor  $\tau$ . Components  $C_1$  and  $C_2$  are only merged if

$$\min [\omega_{\max}(C_1) + \tau(C_1), \omega_{\max}(C_2) + \tau(C_2)] \geq \omega_{\max}(C_F). \quad (7)$$

In practice  $\tau$  is defined as  $\tau(C) = s/|C|$  where  $s$  plays the role of a scale parameter. The above definition presents a few problems. First,  $\tau$  (i.e.  $s$ ) is a global parameter and experiments show that clusters with different sizes and densities might not be recovered with this approach (Figure 10a). Second, there is not an easy rule to fix  $\tau$  or  $s$  and, although it can be related with a scale parameter, there is no way to predict which specific value is best suited to a particular problem.

The exploration of similar ideas, while bearing in mind their shortcomings, leads us to a new clustering method.

## 2.2. Proximal Meaningful Forest

First, let us observe that the edge length distribution of an MST of a configuration of clustered points differs significantly from that of an unclustered point set (Figure 4). As a general idea, by knowing how to model unclustered data, one could detect clustered data by measuring some kind of dissimilarity between both.

In the a contrario framework, there is no need to characterize the elements one wishes to detect but contrarily the elements one wishes to avoid detecting

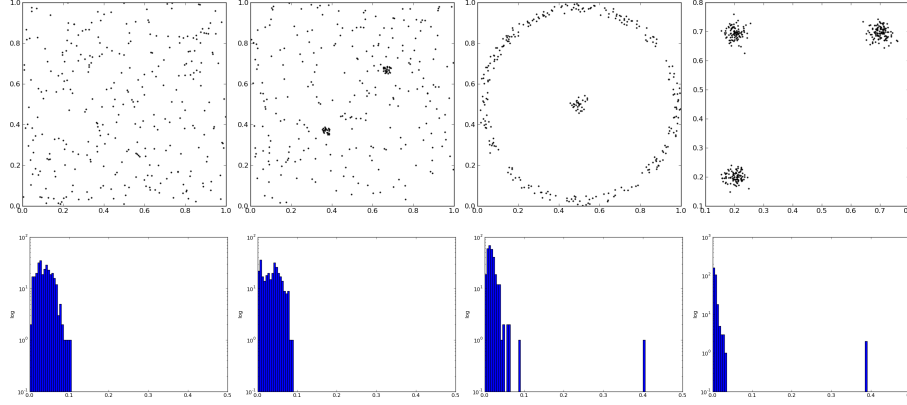


FIG 4. Histograms (in logarithmic scale) of MST edge lengths from different points configurations. The non-clustered case (first column) differs from the other cases. Notice that clustered configurations also differ from each other.

(i.e. unclustered data). To achieve such characterization we only need two elements: (1) a naive model and (2) a measurement made on structures to be potentially detected. The naive model is a probabilistic model that describes typical configurations where no structure is present and we will describe it in detail in Section 2.3. We now focus on establishing a coherent measurement for MST-based clustering.

Concretely, we are looking to evaluate the probability of occurrence, under the background model (i.e. unclustered data), of a random set  $\mathcal{C}$  which exhibits the characteristics of a given observed set  $C$ . Both sets have the same cardinality, i.e.  $|E(C)| = |E(\mathcal{C})| = K$ .

The general principle has been previously explored. In 1983, following the same rationale Hoffman and Jain [17] proposed a similar idea: to perform a test of randomness. They built a null hypothesis using the edge length distribution of the MST and they performed a single test analyzing whether the whole dataset belongs to the random model or not by computing the difference between the theoretical and the empirical CDF. Jain et al. [19] further refined this work, by using heuristic computations to separate the dataset into two or more subsets which were then tested using a two sample test statistic. Barzily et al. recently continued this line of work [4]. This approach introduces a bias towards the detection of compact (i.e. non-elongated) and equally sized clusters [4]. From the perspective chosen in this work these characteristics can be seen as shortcomings, and thus we build a new method.

**Definition 2.** Let  $\mathcal{P}$  be a partition of  $\mathbb{R}$ . We define the equivalence relation  $\sim$  where  $\omega_1 \sim \omega_2$  if and only if  $\exists P \in \mathcal{P}$  such that  $\omega_1 \in P$  and  $\omega_2 \in P$ .

We denote by  $e_i$  the  $i$ -th edge of  $E(C)$  and by  $\gamma_i$  the  $i$ -th edge of  $E(\mathcal{C})$ . Inspired by Equation 7, which proved successful as a decision rule to detect



clusters, and associating it with Equations 5 and 6, we compute

$$\Pr\left(\omega_{\max}(\mathcal{C}) < \omega_{\max}(C) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_F)\right). \quad (8)$$

**Definition 3.** Let  $C \in X$  be a component of the single-link hierarchy  $\mathcal{T}$  induced by the minimum spanning tree  $T = (X, E_T)$  such that  $|C| > 1$ . We define the probability of false alarms (PFA) of  $C$  as

$$\text{PFA}(C) \stackrel{\text{def}}{=} \Pr\left(\omega_{\max}(\mathcal{C}) < \omega_{\max}(C) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_F)\right). \quad (9)$$

The constraint  $|C| > 1$  is needed since  $E(C) = \emptyset$  when  $|C| = 1$ . Note that sets consisting of a single node must certainly not be detected. Conceptually, even when they are isolated, they constitute an outlier and not a cluster. We simply do not test such sets.

To detect unlikely dense subgraphs, a threshold is necessary on the PFA. In the classical *a contrario* framework, a new quantity is introduced: the Number of False Alarms (NFA), i.e. the product of the PFA by the number of tested candidate clusters. The NFA has a more intuitive meaning than the PFA, since it is an upper bound on the expectation of the number of false detections [10]. The threshold is then easily set on the NFA.

**Definition 4** (Number of false alarms). We define the number of false alarms (NFA) of  $C$  as

$$\text{NFA}(C) \stackrel{\text{def}}{=} (|X| - 1) \cdot \text{PFA}(C). \quad (10)$$

Notice that, by definition,  $|X| - 1$  is the number of non-singleton sets in the single-link hierarchy.

**Definition 5** (Meaningful component). A component  $C$  is  $\varepsilon$ -meaningful if

$$\text{NFA}(C) < \varepsilon. \quad (11)$$

In the following, it is important to notice a fact about the single-link hierarchy. The components are mainly determined by the sorted sequence of the edges from the original graph; this follows directly from Kruskal's algorithm [9]. However, the components are independent of the differences between the edges in that sorted sequence: only the order matters and not the actual weights of the edges.

We will now state a simple proposition that motivates Definition 4.

**Proposition 1.** The expected number of  $\varepsilon$ -meaningful clusters in a random single-link hierarchy (i.e. issued from the background model) is lower than  $\varepsilon$ .

The proof is given in Appendix A in the light of the discussion in the next section.

### 2.3. The background model

The distribution  $\Pr\left(\omega_{\max}(\mathcal{C}) < \omega_{\max}(C) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_F)\right)$  is not known a priori. Moreover, up to our knowledge there is no analytical expression for the cumulative edge distribution under  $\mathcal{H}_0$  for the MST [17].

We estimate this distribution by performing Monte Carlo simulations of the background process. However this estimation would involve extremely high computational costs.

We assume that the edge lengths in the cluster, given that  $\omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_F)$ , are mutually conditionally independent and identically distributed. Let  $\Omega$  be a random variable with this common distribution:

$$F_{\Omega}(\omega, \omega_{\max}(C_F)) = \Pr\left(\Omega \leq \omega \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_F)\right). \quad (12)$$

Then, using rank statistics:

$$\begin{aligned} \Pr\left(\omega_{\max}(\mathcal{C}) \leq \omega_{\max}(C) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_F)\right) \\ = F_{\Omega}(\omega_{\max}(C), \omega_{\max}(C_F))^K. \end{aligned} \quad (13)$$

The distribution of  $\Omega$  is much easier to compute than the one of  $\omega_{\max}(\mathcal{C})$  given that  $\omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_F)$  as it requires fewer Monte Carlo simulations and thus we use it as in Equation 13 to compute the PFA.

Now, the independence hypothesis is clearly false even for i.i.d. graph vertices, because of the ordering structure and edge selection induced by the MST. However, the conditional dependence may be weak enough to make the model suitable in practice. This explains why naive classifiers, such as the Naive Bayes classifier, despite their simplicity, can often outperform more sophisticated classification methods [16]. Therefore we only assume conditional independence in our model. Still, the suitability of this model has to be checked; a series of experiments shows that no clusters are detected in non-clustered, unstructured data (for instance, sets of independent, uniformly distributed vertices). On the other side of the problem, meaningful clusters always constitute a clear violation of this naive independence assumption.

The estimation process is depicted in Algorithm 1. Classically, one defines a point process and a sampling window. Hoffman and Jain [17] point out that the sampling window for the background point process is usually unknown for a given dataset. They use the convex hull arguing that it is the maximum likelihood estimator of the true sampling window for uniformly distributed two-dimensional data. In the experiments from Section 3, we simply use the minimum hiper-rectangle that contains the whole dataset as the sampling window. However, there are problems whose intrinsic characteristics allow to define other background processes that do not involve a sampling window.

---

**Algorithm 1** Compute  $\Pr \left( \omega_{\max}(\mathcal{C}) < \omega_{\max}(C) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_F) \right)$  for a set of  $N$  points by  $Q$  Monte Carlo simulations.

---

```

for all  $q$  such that  $1 \leq q \leq Q$  do
   $X \leftarrow$  draw  $N$  points from the background point process.
  build the single-link hierarchy  $\mathcal{T}_q$  from the MST of  $X$ .
end for
compute a conditional histogram from the set  $\{\mathcal{T}_q\}_{q=1\dots Q}$ 

```

---

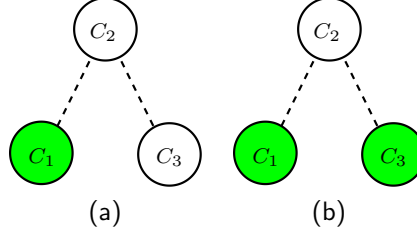


FIG 5. *Example of collateral elimination. Three components  $C_1, C_2, C_3$  such that  $C_1 \subset C_2$ ,  $C_3 \subset C_2$  and  $\text{NFA}(C_1) < \text{NFA}(C_2) < \text{NFA}(C_3) < \varepsilon$ . (a) The classical maximality rule only selects  $C_1$  as a maximal component. (b) The scheme in Algorithm 2 selects  $C_1$  and  $C_3$ .*

#### 2.4. Eliminating redundancy

While each meaningful cluster is relevant by itself, the whole set of meaningful components exhibits, in general, high redundancy: a meaningful component  $C_1$  can contain another meaningful component  $C_2$  [5]. This question can be answered by comparing  $\text{NFA}(C_1)$  and  $\text{NFA}(C_2)$  using Definition 5. The group with the smallest NFA must of course be preferred. Classically, the following rule

```

for all  $\varepsilon$ -meaningful clusters  $C_1, C_2$  do
  if  $C_2 \subset C_1 \vee C_1 \subset C_2$  then
    eliminate  $\arg \max (\text{NFA}(C_1), \text{NFA}(C_2))$ 
  end if
end for

```

would have been used to perform the pruning of the set of meaningful components. Unfortunately, it leads to a phenomenon described in [5], where it was called collateral elimination. A very meaningful component can hide another meaningful sibling, as illustrated in Figure 5.

The single-link hierarchy offers an alternative scheme to prune the redundant set of meaningful components, profiting from the inclusion properties of the dendrogram structure. It is non-other than the exclusion principle, defined first by Desolneux et al. [10], which states that

Let  $A$  and  $B$  be groups obtained by the same gestalt law. Then no point  $x$  is allowed to belong to both  $A$  and  $B$ . In other words each point must either belong to  $A$  or to  $B$ .

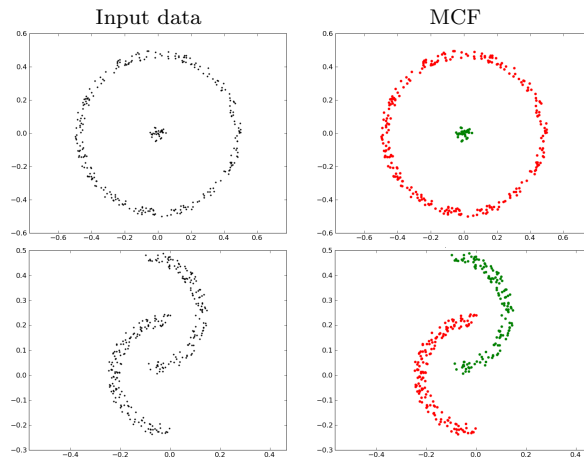


FIG 6. The meaningful clustered forest correctly describes the points organization, even when clusters have arbitrary shapes.

A simple scheme for applying the exclusion principle is shown in Algorithm 2.

Since we are choosing the components that are more in accordance with the proximity gestalt, we call the resulting components Proximal Meaningful Components (PMC). Then, we say that the set of all proximal meaningful components is a Meaningful Clustered Forest (MCF).

---

**Algorithm 2** Eliminate redundant components from the set  $\mathcal{M}$  of meaningful components.

---

```

1:  $\mathcal{F} \leftarrow \emptyset$ 
2: while  $\mathcal{M} \neq \emptyset$  do
3:    $C_{\min} \leftarrow \arg \min_{C \in \mathcal{M}} \text{NFA}(C)$ 
4:   eliminate  $C_{\min}$  from  $\mathcal{M}$ 
5:   eliminate all components  $C$  from  $\mathcal{M}$  such that  $C \subset C_{\min}$ 
6:   eliminate all components  $C$  from  $\mathcal{M}$  such that  $C_{\min} \subset C$ 
7:   add  $C_{\min}$  to  $\mathcal{F}$ 
8: end while
9:  $\mathcal{M} \leftarrow \mathcal{F}$ 

```

---

### 3. Experiments on Synthetic examples

As a sanity check, we start by testing our method with simple examples. Figure 6 present clusters which are well but not linearly separated. The meaningful clustered forest describes correctly the structure of the data.

Figure 7 shows an example of cluster detection in a dataset overwhelmed by outliers. The data consists of 950 points uniformly distributed in the unit square, and 50 points manually added around the positions  $(0.4, 0.4)$  and  $(0.7, 0.7)$ . The figure shows the result of a numerical method involving the above NFA. The

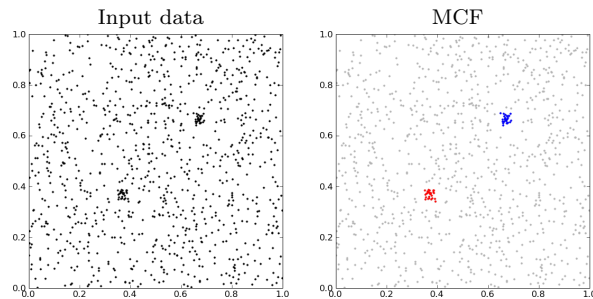


FIG 7. Similar experiment as performed by Cao et al. in Figure 2 [6]. Clustering of twice 25 points around  $(0.4, 0.4)$  and  $(0.7, 0.7)$  surrounded by 950 i.i.d. points, uniformly distributed in the unit square. Exactly two proximal meaningful components are detected.

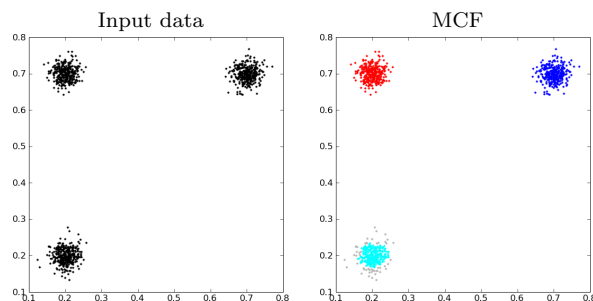


FIG 8. Clusters are correctly recovered in the mixture of three Gaussians. However some points are detected as noise (depicted in gray) in the tails.

background distribution is chosen to be uniform in  $[0, 1]^2$ . Both visible clusters are found and their NFAs are respectively  $10^{-15}$  and  $10^{-8}$ . Such low numbers can barely be the result of chance.

The case of mixture of Gaussians, shown in Figure 8, provides an interesting example. On the tails, points are obviously sparser and the distance to neighboring points grows. Since we are looking for tight components, the tail might be discarded, depending on the Gaussian variance.

The example in Figure 9 consists of a very complex scene, composed of clusters with different densities, shapes and sizes. Proximal components (i.e. we avoid testing  $\text{NFA} < \varepsilon$ ) are displayed. Even when no decision about the statistical significance is made, the recovered clusters describe, in general, the scene accurately. Some oversplitting can be detected in proximal components. When a decision is made and only meaningful components are kept, we realize that the oversplit figures are not meaningful. As a sanity check, in Figure 9e we plot some of the detected structures superimposed to a realization of the background noise model. The input data in Figure 9a contains 764 points and for a given shape in it, with  $W$  points, we plot the shape and  $764 - W$  points drawn from the background model. Among proximal components, the meaningful ones can

be clearly perceived while non-meaningful ones are unnoticed.

Our results are compared with Felzenszwalb and Huttenlocher' algorithm (denoted by FH in the following) and with Mean Shift [8, 14]. Mean Shift performs a non-parametric density estimation (using sliding windows) and finds its local maxima. Clusters are determined by what Comaniciu and Meer call "bassins of attraction" [8]: points are assigned to a local maximum following an ascendent path along the density gradient <sup>1</sup>.

Figure 10 present an experiment were FH and Mean Shift are used, respectively, to cluster the dataset in Figure 9a. Different runs were performed, by varying the kernel/scale size. Clearly, results are suboptimal in both cases. Both algorithms share the same main disadvantage: a global scale must be chosen a priori. Such a strategy is unable to cope with clusters of different densities and spatial sizes. Choosing a fine scale causes to correctly detect dense clusters at the price of oversplitting less denser ones. On the contrary, a coarser scale corrects the oversplitting of less denser clusters but introduces undersplitting for the denser ones.

#### 4. Handling MST Instability

A seemingly obvious but interesting phenomenon occurs when noise is added to clustered data. Suppose we have data with two well separated clusters. In the absence of noise, it exists an MST edge linking both clusters. If noise is added to the data, the edge would probably disappear and be replaced by a sequence of edges. The length of the original linking edge is larger than the length of the endpoints of the sequence. The direct consequence is an increase in the NFA of the two clusters. Depending on the magnitude of that increase, the clusters might potentially be split into several proximal meaningful components. See Figure 13.

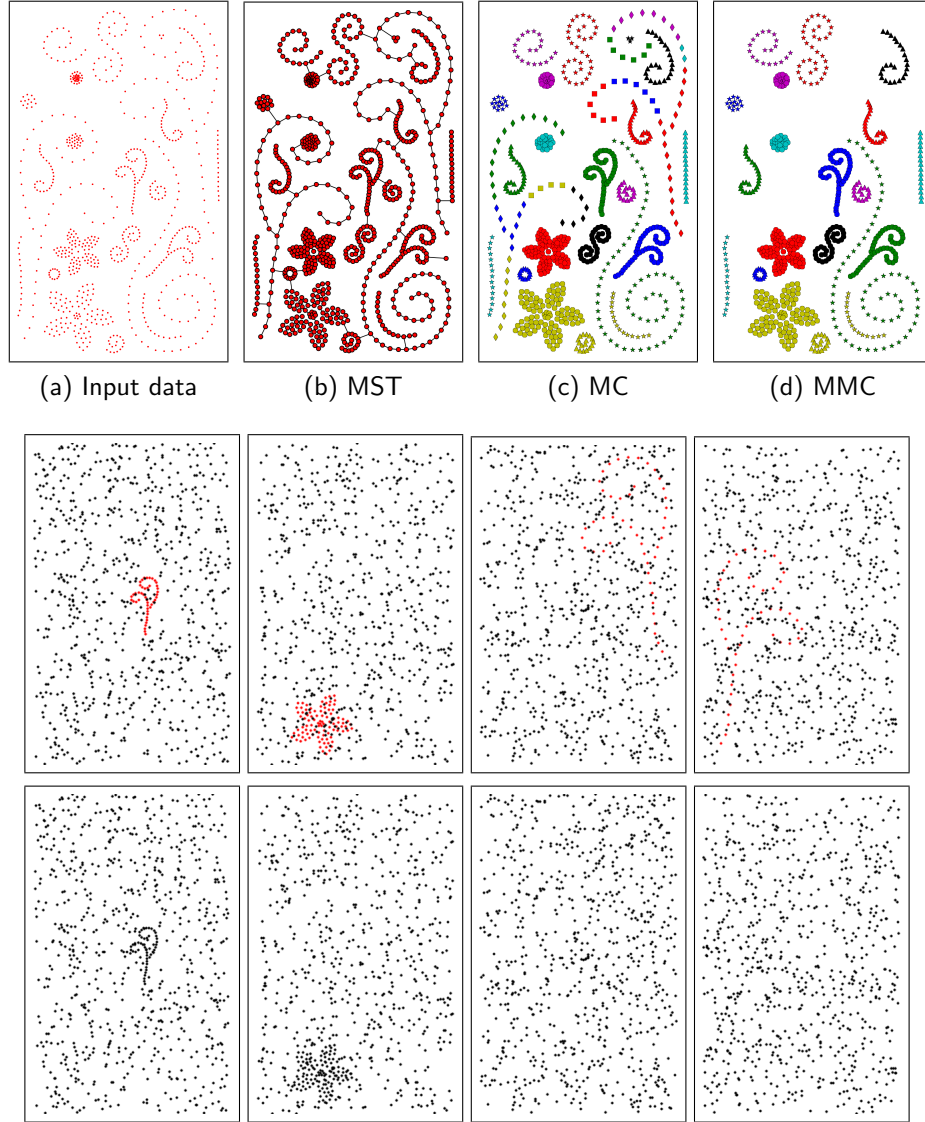
In short terms, noise affects the ideal topology of the MST. The oversplitting phenomenon can be corrected by iterating the following steps:

1. detecting the meaningful clustered forest,
2. add the union of points in the meaningful clustered forest to a new input dataset,
3. remove the points in the meaningful clustered forest and replace them with noise,
4. iterate until convergence,
5. re-detect the meaningful clustered forest on the new noise-free dataset built along all iterations.

The MST of the set formed by merging the meaningful clustered forests from all iterations has the right topology. In other words this MST resembles the MST of the original data without noise. Then, detection of meaningful clustered forest

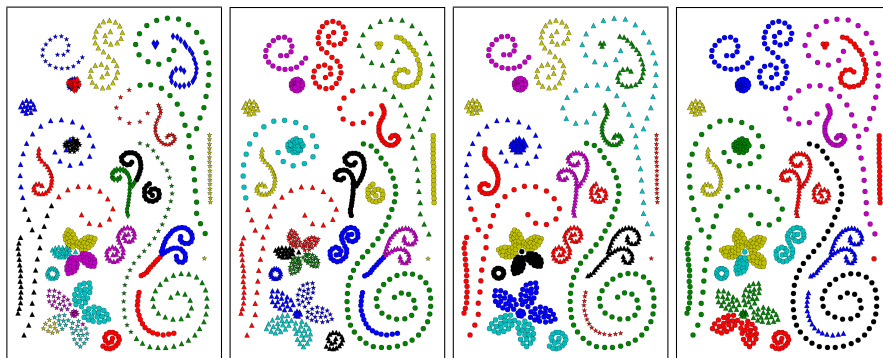
---

<sup>1</sup>code available at <http://www.mathworks.com/matlabcentral/fileexchange/10161-mean-shift-clustering>

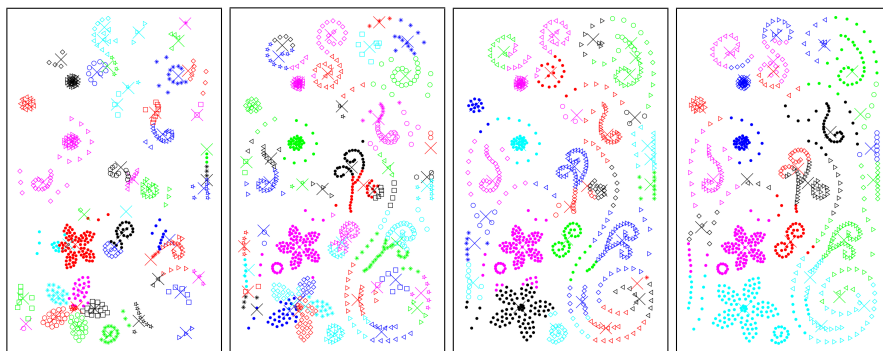


(e) Shapes drawn against noise. Shapes are respectively plotted in red and in black on the top and bottom rows.

FIG 9. In this example, maximal meaningful components correctly describe the organization of the points configuration. Only small or less denser figures are discarded. Indeed, meaningful components are clearly perceived in noise while non-meaningful components are not.



(a) Felzenszwalb and Huttenlocher' algorithm results at different scales.



(b) Mean Shift results with different values of kernel sizes.

FIG 10. The same points configuration as in Figure 9a. At all scales, both algorithms fail to correctly detect the organization. Under and oversplitting occur in all cases.

can be performed without major trouble. We say that these detections form a stabilized meaningful clustered forest.

The above method implicitly contains a key issue in step 3. Detected points must be replaced with others which have a completely different distribution (i.e. the background distribution) but must occupy the same space. Figure 11 contains an example of the need for such a strong requirement. Pieces of background data might become “disconnected, or to be precise connected by artificially created new edges. In one dimension, these holes are easily contracted, but when the dimensionality increases the contracting scheme gains more and more degrees of freedom.

This noise filling procedure can be achieved by using the Voronoi diagram [2] of the original point set. In the Voronoi diagram, each point lies on a different cell. To remove a point amounts to emptying a cell. Then the set of empty cells can be used as a sampling window to draw points from the background model. Notice that this procedure is actually generic since the Voronoi tessellation can be generalized to higher dimensional metric spaces [2].



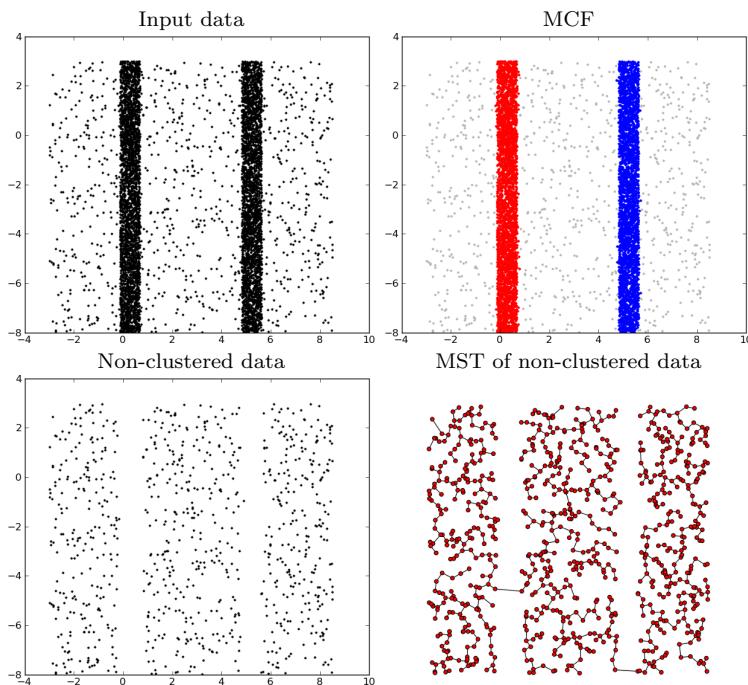


FIG 11. Removing PMC can generate artifacts, *i.e.* holes, in the remaining data. These holes might create edges in the MST of the non-clustered data, that certainly violate the background model.

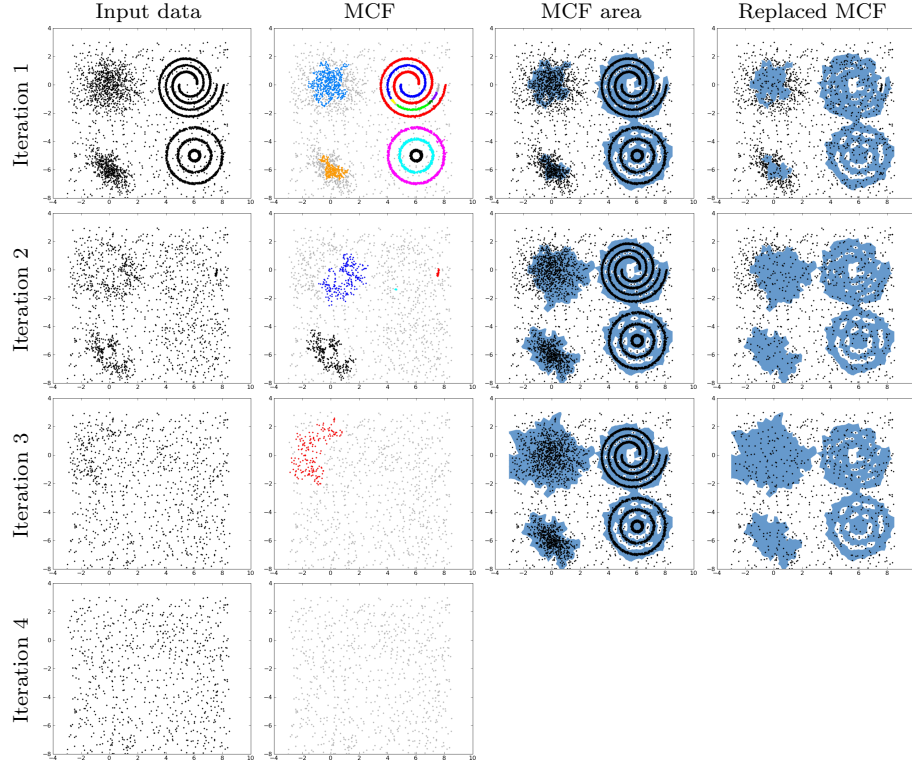
The process simulates replacing detected components with noise from the background model. Due to the same nature of the Voronoi diagram, the process is not perfect: in the final iteration, the resulting point set is quasi but not exactly distributed according to the background model. A small bias is introduced, causing a few spurious detections in the MCF. To correct this issue it suffices to set  $\varepsilon = 10^{-2}$ , as these detections have NFAs slightly lower than one and actual detections have really low NFAs. Of course this new thresholding could be avoided if a more accurate filing procedure was used.

Algorithm 3 illustrates steps 1 to 4 of the correcting method. An example is shown in Figure 12, where four iterations are required until convergence.

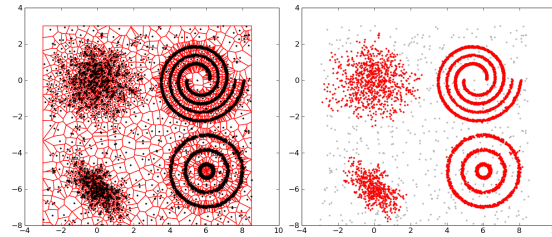
Figure 13 shows a second example of the stabilization process, followed by the detection of the stabilized meaningful clustered forest. The NFAs of the detected components are also included. The very low attained NFAs, account for the success of the procedure.

## 5. The Masking Challenge

In 2009, Jain [18] stated that no available algorithm could cluster the dataset in Figure 14b and obtain the result in Figure 14c. The dataset is interesting



(a) In each iteration, the MCF is detected and the cells on the Voronoi diagram corresponding to points in the MCF are emptied and filled with points distributed according to the background model. In the fourth iteration, no MCF is detected and thus the algorithm stops.



(b) Left, original Voronoi diagram. Right, resulting non-background points in red.

FIG 12. Iteratively detecting the MCF and replacing it with points from the background model, converges and separates background from non-background data.

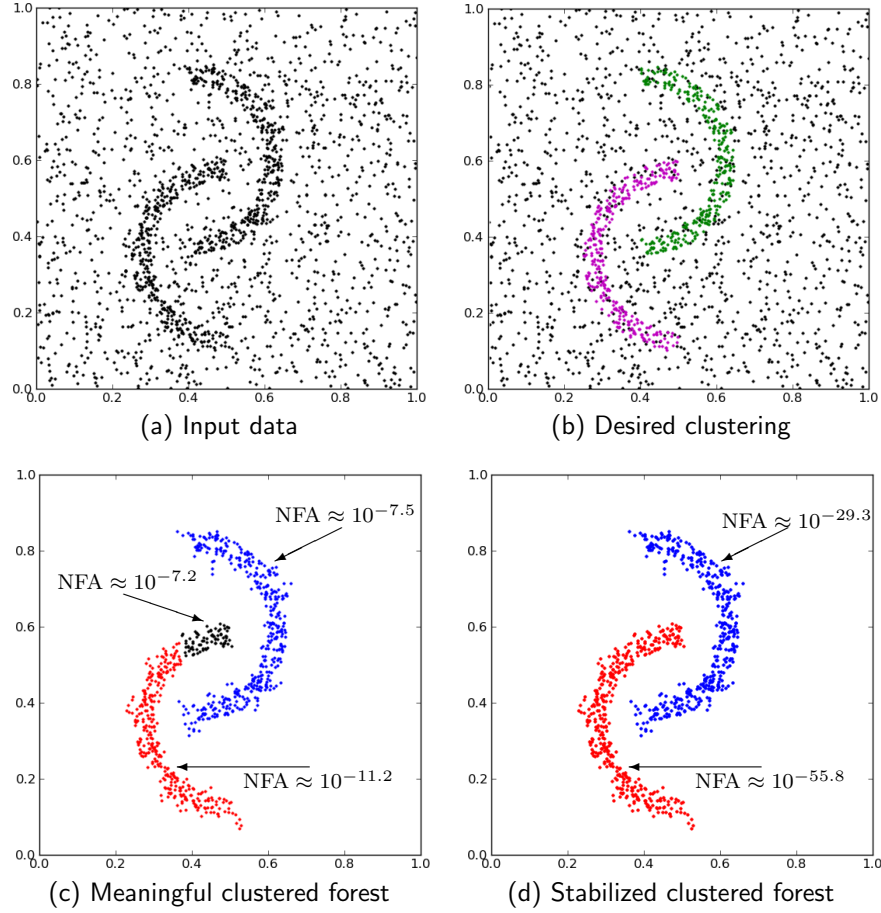


FIG 13. Noise might affect the stability of the meaningful clustered forest, causing to oversplit the clusters. Algorithm 3 converges in two iterations. Then, one can detect the meaningful clustered forest among non-background points, yielding a stabilized meaningful clustered forest.

---

**Algorithm 3** Stabilize point set  $X$  returning the set  $\mathcal{F}$  of non-background points.

---

```

1:  $\mathcal{F} \leftarrow \emptyset$ 
2:  $\mathcal{V} \leftarrow$  cells from Voronoi diagram of point set  $X$  intersected with the minimum rectangle
   enclosing  $X$ .
3:  $X' \leftarrow X$ 
4:  $\mathcal{M} \leftarrow$  meaningful clustered forest of  $X'$ 
5: while  $\mathcal{M} \neq \emptyset$  do
6:    $\mathcal{V}' \leftarrow \emptyset$ 
7:    $X' \leftarrow \emptyset$ 
8:   for all  $C \in \mathcal{M}$  do
9:     for all  $\mathbf{p} \in C$  do
10:      add  $V \in \mathcal{V}$  to  $\mathcal{V}'$  such that  $\mathbf{p} \in V$ .
11:      if  $\mathbf{p} \in X$  then
12:        add  $\mathbf{p}$  to  $X'$ .
13:        add  $\mathbf{p}$  to  $\mathcal{F}$ .
14:      end if
15:    end for
16:   end for
17:    $a \leftarrow \sum_{V \in \mathcal{V}} \text{area}(V)$ 
18:    $a_{\mathcal{M}} \leftarrow \sum_{V \in \mathcal{V}'} \text{area}(V)$ 
19:    $n_{\mathcal{M}} \leftarrow \sum_{C \in \mathcal{M}} |C|$ 
20:    $n \leftarrow a_{\mathcal{M}} \cdot (|X| - n_{\mathcal{M}}) / (a - a_{\mathcal{M}})$ 
21:    $B \leftarrow$  draw  $n$  points  $\mathbf{q}_i$ ,  $1 \leq i \leq n$ , from the background model such that  $(\exists V \in \mathcal{V}') q_i \in V$ .
22:    $X' \leftarrow X' \cup B$ 
23:    $\mathcal{M} \leftarrow$  meaningful clustered forest of  $X'$ 
24: end while

```

---

because it brings to light a new phenomenon: a cluster with many points can “dominate” the scene and hide other clusters that could be meaningful.

A similar behavior occurs in vanishing point detection, as pointed out by Almansa et al. [1]. A vanishing point is a point in an image to which parallel line segments not frontoparallel appear to converge; in some sense one can regard this point as a collection of such parallel line segments. Sometimes this procedure will still miss some weak vanishing points which are “masked” by stronger vanishing points composed of much more segments. These may not be perceived at first sight, but only if we manage to unmask them by getting rid of the “clutter” in one way or another. Almansa et al. propose to eliminate these detected vanishing points and look for new vanishing points among the remaining line segments.

In our case, this very same approach cannot be followed. Masked clusters are not completely undetected but partially detected. Removing such cluster parts and re-detecting would cause oversegmentation. We propose instead to only remove the most meaningful proximal component and then iterate. The process ends when the masking phenomenon disappears, that is:

- when there are no unclustered points, or
- no MCF is detected.

Algorithm 4 shows a detail of this unmasking scheme. Summarizing, first non-background points are detected using the stabilization process in Algorithm 3 and then the unmasking process takes place.

The detection of unmasked MMCs in Figure 14d correct all masking issues. Moreover, they are extremely similar to the desired clustering in Figure 14c. The difference is that clusters absorb background points that are within of near them. Indeed, these background points are statistically indistinguishable from the points from the cluster that absorbs them.

---

**Algorithm 4** Compute the unmasked meaningful clustered forest  $\mathcal{U}$  from the point set  $\mathcal{F}$  of non-background points.

---

```

1:  $\mathcal{U} \leftarrow \emptyset$ 
2: while  $\mathcal{M} \neq \emptyset$  do
3:    $\mathcal{M} \leftarrow$  stabilized meaningful clustered forest of  $\mathcal{F}$ 
4:   if  $|\mathcal{F}| = \sum_{C \in \mathcal{M}} |C|$  then
5:      $\forall C \in \mathcal{M}$ , add  $C$  to  $\mathcal{U}$ .
6:   else
7:      $C_{\min} \leftarrow \arg \min_{C \in \mathcal{M}} \text{NFA}(C)$ 
8:     for all  $\mathbf{p} \in C_{\min}$  do
9:       remove  $p$  from  $\mathcal{F}$ 
10:    end for
11:    add  $C_{\min}$  to  $\mathcal{U}$ .
12:  end if
13: end while

```

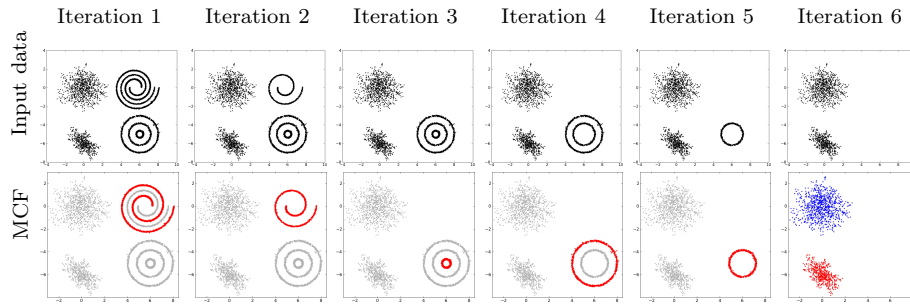
---

From a total number of 7000 points in Figure 14b, the outer spiral (in orange in in Figure 14c) has 2514 points, i.e. almost 36% of the points. The detection of the unmasked MCF in Figure 14d correct all masking issues. Moreover, they are extremely similar to the desired clustering in Figure 14c. The difference is that clusters absorb background points that are within of near them. Indeed, these background points are statistically indistinguishable from the points from the cluster that absorbs them.

## 6. Three-dimensional point clouds

We tested the proposed algorithm with three-dimensional point clouds. We put two point clouds in the same scene at different positions, thus building two scenes in Figures 15 and 16. In both cases uniformly distributed noise was artificially added. The skeleton hand and the bunny are formed by 3274 and by 3595, respectively. In Figure 15, 3031 noise points were added to total 9900 points. In Figure 15, 7031 noise points were added to total 13900 points and both shapes were positioned closer to each other and in such a way that no linear separation exist between them. In both cases the result is correct

In Figure 15, the MCF is oversplit but the stabilization process discussed in Section 4 corrects the issue. In Figure 15, although the same phenomenon is possible, it does not occur in this realization of the noise process.



(a) In each iteration, the MCF is detected and the most meaningful component is removed from the dataset. In the sixth iteration, all points are clustered and thus the algorithm stops.

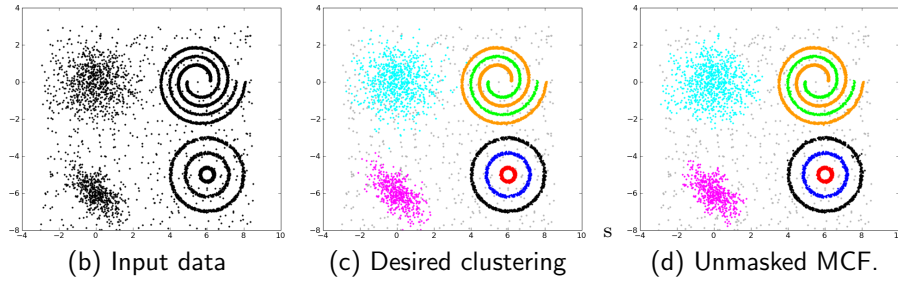


FIG 14. Iteratively detecting the MCF and removing from the dataset its most meaningful component, converges and corrects the masking phenomenon. The detected MCF is extremely similar to the desired clustering. The difference is that clusters absorb background points that are within of near them.

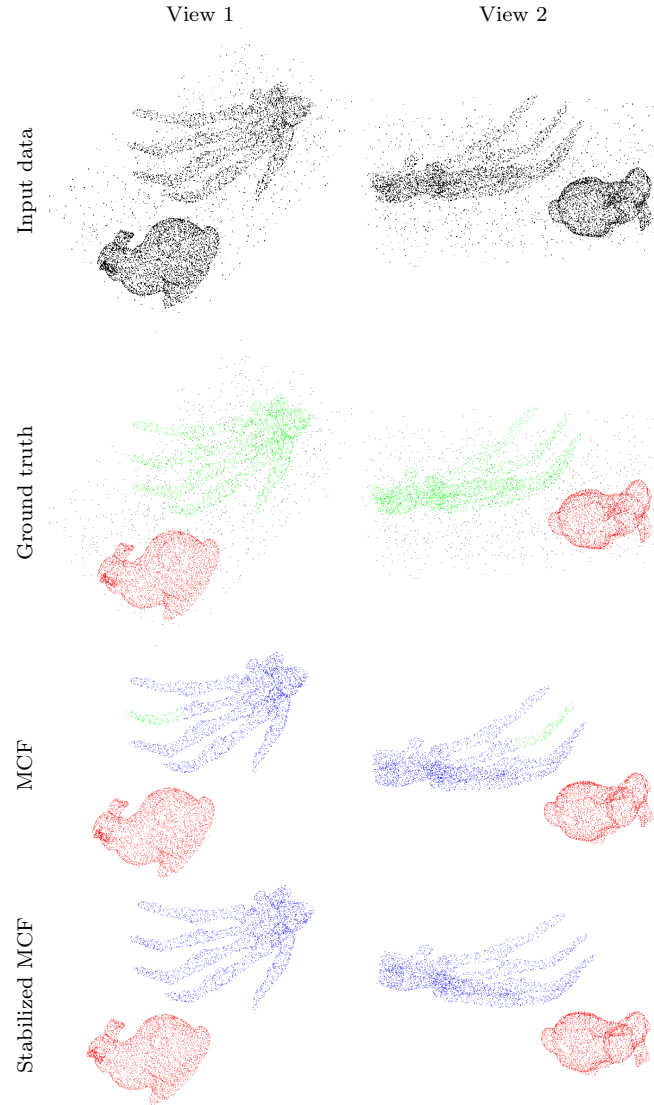


FIG 15. Two point clouds with artificially added noise. In this case, noise perturbed the MCF (see the finger of the skeleton hand). This effect is corrected by the stabilization process.

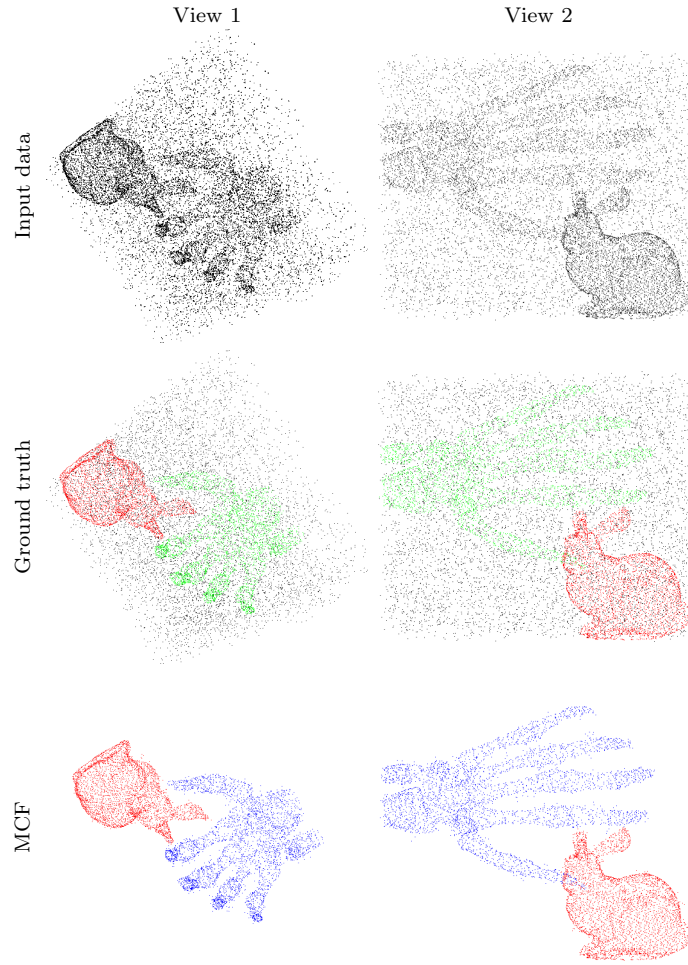


FIG 16. Two point clouds with artificially added noise. Both shapes are close to each other and are not linearly separable. The result of the stabilization process is omitted as detections do not change.



## 7. Final Remarks

In this work we propose a new clustering method that can be regarded as a numerical method to compute the proximity gestalt. The method relies on analyzing edge distances in the MST of the dataset. The direct consequence is that our approach is fully parametric on the chosen distance.

The proposed method presents several novelties over other MST-based formulations. Some formulations have preference for compact clusters as they extract their clustering detection rule from characteristics that are not intrinsic to the MST. Our method only focuses on the length of the MST edges; hence, it does not present such preference. Other formulations analyze the data at a fixed local scale, thus introducing a new method parameter. We have shown through examples that these local methods can fail when the input data has clusters with different sizes and densities. In these same examples, our method performs well without the need of introducing any extra parameter.

The method is also automatic, in the sense that only a single parameter is left to the user. This parameter has an intuitive interpretation as it controls the expected number of false detections. Moreover, setting it to 1 is sufficient in practice.

Robustness to noise is an additional but essential feature of the method. Indeed, we have shown that the iterative application of our method can be used to treat noisy data, producing quality results.

We also studied the masking phenomenon in which a highly populated and salient cluster dominates the scene and inhibits the detection of less-populated, but still salient, clusters. The proposed method can be iteratively used to avoid such inhibitions from happening, yielding promising results.

As future work, it would be interesting to study the MST edge distribution under different point distributions. From the theoretical point of view, it can bring light to the method correctness. In practice, it would allow to replace the simulated background models by their analytical counterparts.

## Appendix A: Proof of Proposition 1

The proof relies on the following classical lemma.

**Lemma 1.** *Let  $X$  be a real random variable and let  $F(x) = P(X \leq x)$  be the cumulative density function of  $X$ . Then for all  $t \in (0, 1)$ ,*

$$\Pr(F(X) < t) \leq t \tag{14}$$

**Proposition 2.** *The expected number of  $\varepsilon$ -meaningful clusters in a random single-link hierarchy (i.e. issued from the background model) is lower than  $\varepsilon$ .*

*Proof.* We follow the scheme of Proposition 1 from the work by Cao et al. [6]. Let  $\mathcal{T}$  be random single-link hierarchy. For brevity let  $M = |X| - 1$ . Let  $Z_i$  be

a binary random variable equal to 1 if the random cluster  $C_i \in \mathcal{T}$  is meaningful and 0 else.

Let us denote by  $\mathbb{E}(X)$  the expectation of a random variable  $X$  in the a contrario model. We then have

$$\mathbb{E} \left( \sum_{i=1}^M Z_i \right) = \mathbb{E} \left( \mathbb{E} \left( \sum_{i=1}^M Z_i \mid M \right) \right). \quad (15)$$

Let  $Y_i$  be a binary random variable equal to 1 if

$$M \cdot \Pr \left( \omega_{\max}(\mathcal{C}) < \omega_{\max}(C_i) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_{iF}) \right) < \varepsilon \quad (16)$$

and 0 else. Of course,  $M$  is independent from the sets in  $\mathcal{T}$ . Thus, conditionally to  $M = m$ , the law of  $\sum_{i=1}^M Z_i$  is the law of  $\sum_{i=1}^m Y_i$ . Let us reprise Equation 13 on p. 10,

$$\begin{aligned} \Pr \left( \omega_{\max}(\mathcal{C}) < \omega_{\max}(C_i) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_{iF}) \right) \\ = F_{\Omega} \left( \omega_{\max}(C_i), \omega_{\max}(C_{iF}) \right)^{K_i}. \end{aligned} \quad (17)$$

By linearity of expectation,

$$\mathbb{E} \left( \sum_{i=1}^M Z_i \mid M = m \right) = \mathbb{E} \left( \sum_{i=1}^m Y_i \right) = \sum_{i=1}^m \mathbb{E}(Y_i). \quad (18)$$

Let us denote  $F_{\Omega} \left( \omega_{\max}(C_i), \omega_{\max}(C_{iF}) \right)$  by  $\Pr(C_i)$ . Since  $Y_i$  is a Bernoulli variable,

$$\begin{aligned} \mathbb{E}(Y_i) &= \Pr(Y_i = 1) = \Pr \left( M \cdot \Pr(C_i)^{K_i} < \varepsilon \right) \\ &= \sum_{k=0}^{\infty} \Pr \left( M \cdot \Pr(C_i)^{K_i} < \varepsilon \mid K_i = k \right) \cdot \Pr(K_i = k). \end{aligned} \quad (19)$$

We have assumed that  $K_i$  is independent from  $\Pr(C_i)$ . Thus, conditionally to  $K_i = k$ ,  $\Pr(C_i)^{K_i} = \Pr(C_i)^k$ . We have

$$\begin{aligned} &\Pr \left( m \cdot \Pr \left( \omega_{\max}(\mathcal{C}) < \omega_{\max}(C_i) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_{iF}) \right) < \varepsilon \right) \\ &= \Pr \left( m \cdot F_{\Omega} \left( \omega_{\max}(C_i), \omega_{\max}(C_{iF}) \right)^k < \varepsilon \right) \\ &= \Pr \left( m \cdot \Pr \left( \Omega < \omega_{\max}(C_i) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_{iF}) \right)^k < \varepsilon \right) \\ &= \Pr \left( \Pr \left( \Omega < \max_{e \in C_i} \omega(e) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_{iF}) \right) < \left( \frac{\varepsilon}{m} \right)^{1/k} \right) \end{aligned}$$

$$\begin{aligned}
&= \Pr \left( \max_{e \in E(C_i)} \Pr \left( \Omega < \omega(e) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_{iF}) \right) < \left( \frac{\varepsilon}{m} \right)^{1/k} \right) \\
&= \prod_{e \in E(C_i)} \Pr \left( \Pr \left( \Omega < \omega(e) \mid \omega_{\max}(\mathcal{C}_F) \sim \omega_{\max}(C_{iF}) \right) < \left( \frac{\varepsilon}{m} \right)^{1/k} \right). \quad (20)
\end{aligned}$$

The last equality follows from the conditional independence assumption. Now, using Lemma 1 and bearing in mind that the number of edges in  $C_i$  is  $K_i = k$ , yields

$$\Pr \left( m \cdot \Pr(C_i)^k < \varepsilon \right) \leq \prod_{j=1}^k \left( \frac{\varepsilon}{m} \right)^{1/k} = \frac{\varepsilon}{m}. \quad (21)$$

We can now use this bound in Equation 19:

$$\begin{aligned}
\mathbb{E}(Y_i) &= \sum_{k=0}^{\infty} \Pr \left( M \cdot \Pr(C_i)^{K_i} < \varepsilon \mid K_i = k \right) \cdot \Pr(K_i = k) \\
&\leq \frac{\varepsilon}{m} \sum_{k=0}^{\infty} \Pr(K_i = k) = \frac{\varepsilon}{m}. \quad (22)
\end{aligned}$$

Hence,

$$\mathbb{E} \left( \sum_{i=1}^M Z_i \mid M = m \right) = \sum_{i=1}^m \mathbb{E}(Y_i) \leq \varepsilon. \quad (23)$$

This finally implies  $\mathbb{E} \left( \sum_{i=1}^M Z_i \right) \leq \varepsilon$ , what means that the expected number of  $\varepsilon$ -meaningful clusters is less than  $\varepsilon$ .  $\square$

## Acknowledgments

The authors acknowledge financial support by CNES (R&T Echantillonnage Ir-régulier DCT / SI / MO - 2010.001.4673), FREEDOM (ANR07-JCJC-0048-01), Callisto (ANR-09-CORD-003), ECOS Sud U06E01, ARFITEC (07 MATRH) and STIC Amsud (11STIC-01 - MMVPSCV) and the Uruguayan Agency for Research and Innovation (ANII) under grant PR-POS-2008-003.

## References

- [1] A. Almansa, A. Desolneux, and S. Vamech. Vanishing point detection without any a priori information. *Transactions on Pattern Analysis and Machine Intelligence*, 25(4):502–507, 2003.
- [2] F. Aurenhammer. Voronoi diagrams - a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3):345–405, September 1991.
- [3] S. Bandyopadhyay. An automatic shape independent clustering technique. *Pattern Recognition*, 37(1):33–45, January 2004.

- [4] Z. Barzily, Z. Volkovich, B. Akteke Öztürk, and G. W. Weber. On a Minimal Spanning Tree Approach in the Cluster Validation Problem. *Informatica*, 20(2):187–202, 2009.
- [5] F. Cao, J. Delon, A. Desolneux, P. Musé, and F. Sur. A Unified Framework for Detecting Groups and Application to Shape Recognition. *Journal of Mathematical Imaging and Vision*, 27(2):91–119, February 2007.
- [6] F. Cao, P. Musé, and F. Sur. Extracting Meaningful Curves from Images. *J. Math. Imaging Vis.*, 22(2-3):159–181, 2005.
- [7] G. Carlsson and F. Mémoli. Characterization, Stability and Convergence of Hierarchical Clustering methods. *Journal of Machine Learning Research*, 11:1425–1470, 2010.
- [8] D. Comaniciu and P. Meer. Mean shift: a robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):603–619, August 2002.
- [9] T. Cormen, C. Leiserson, R. Rivest, and C. Stein. *Introduction to Algorithms*. McGraw-Hill Science / Engineering / Math, 2nd edition, December 2003.
- [10] A. Desolneux, L. Moisan, and J. M. Morel. *From Gestalt Theory to Image Analysis*, volume 34. Springer-Verlag, 2008.
- [11] M. Dry, D. Navarro, K. Preiss, and M. Lee. The Perceptual Organization of Point Constellations. In *Annual Meeting of the Cognitive Science Society*, 2009.
- [12] P. Felzenszwalb and D. Huttenlocher. Efficient Graph-Based Image Segmentation. *International Journal of Computer Vision*, 59(2):167–181, September 2004.
- [13] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, second edition, 1999.
- [14] K. Fukunaga and L. Hostetler. The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Transactions on Information Theory*, 21(1):32–40, January 1975.
- [15] R. Grompone von Gioi, J. Jakubowicz, J. M. Morel, and G. Randall. LSD: A Fast Line Segment Detector with a False Detection Control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(4):722–732, April 2010.
- [16] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning, Second Edition: Data Mining, Inference, and Prediction*. Springer Series in Statistics. Springer, 2nd edition, February 2009.
- [17] Richard Hoffman and Anil K. Jain. A test of randomness based on the minimal spanning tree. *Pattern Recognition Letters*, 1(3):175–180, March 1983.
- [18] A. K. Jain. Data clustering: 50 years beyond K-means. *Pattern Recognition Letters*, 31(8):651–666, September 2010.
- [19] A. K. Jain, Xiaowei Xu, Tin K. Ho, and Fan Xiao. Uniformity testing using minimal spanning tree. In *International Conference on Pattern Recognition*, pages 281–284. IEEE Computer Society, 2002.
- [20] R. Kannan, S. Vempala, and A. Vetta. On clusterings: Good, bad and

- spectral. *Journal of the ACM*, 51(3):497–515, May 2004.
- [21] J. Kleinberg. An impossibility theorem for clustering. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems*, pages 446–453. MIT Press, 2002.
  - [22] J. Rabin, J. Delon, and Y. Gousseau. A Statistical Approach to the Matching of Local Features. *SIAM Journal on Imaging Sciences*, 2(3):931–958, 2009.
  - [23] M. Tepper, P. Musé, A. Almansa, and M. Mejail. Automatically finding clusters in normalized cuts. *Pattern Recognition*, 44(7):1372–1386, July 2011.
  - [24] M. Tepper, P. Musé, A. Almansa, and M. Mejail. Boruvka Meets Nearest Neighbors. Technical report, Departamento de Computación, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires, Argentina, April 2011.
  - [25] M. Wertheimer. *Laws of organization in perceptual forms*, pages 71–88. Routledge and Kegan Paul, 1938.
  - [26] C. T. Zahn. Graph-Theoretical Methods for Detecting and Describing Gestalt Clusters. *Transactions on Computers*, C-20(1):68–86, 1971.